DIFFUSION NEGATIVE PREFERENCE OPTIMIZATION MADE SIMPLE

Anonymous authors

000

001

002003004

010 011

012

013

014

015

016

017

018

019

021

024

025

026

027 028 029

031

032

033

034

035

037

040

041

042

043

044

046

047

051

052

Paper under double-blind review

ABSTRACT

Classifier-Free Guidance (CFG) improves diffusion sampling by encouraging conditional generations while discouraging unconditional ones. Existing preference alignment methods, however, focus only on positive preference pairs, limiting their ability to actively suppress undesirable outputs. Diffusion Negative Preference Optimization (Diff-NPO) approaches this limitation by introducing a separate negative model trained with inverted labels, allowing it to capture signals for suppressing undesirable generations. However, this design comes with two key drawbacks. First, maintaining two distinct models throughout training and inference substantially increases computational cost, making the approach less practical. Second, at inference time, Diff-NPO relies on weight merging between the positive and negative models, a process that dilutes the learned negative alignment and undermines its effectiveness. To overcome these issues, we introduce Diff-SNPO, a single-network framework that jointly learns from both positive and negative preferences. Our method employs a bounded preference objective to prevent winner-likelihood collapse, ensuring stable optimization. Diff-SNPO delivers strong alignment performance with significantly lower computational overhead, showing that explicit negative preference modeling can be simple, stable, and efficient within a unified diffusion framework. Code will be released.

1 Introduction

Diffusion models (Ho et al., 2020) have become the backbone of modern visual content generation, achieving remarkable fidelity in synthesizing images, videos, and multimodal content (Rombach et al., 2022; Ho et al., 2022; Ruan et al., 2023). However, models trained on vast, uncurated webscale datasets often inherit biases and fail to align with human notions of quality, aesthetics, or safety. Consequently, fine-tuning models with human feedback through preference alignment has become a critical step for bridging the gap between a model's raw capabilities and user intent (Black et al., 2024; Wallace et al., 2024; Li et al., 2024).

At the heart of high-quality diffusion sampling lies *Classifier-Free Guidance (CFG)* (Ho & Salimans, 2021), a technique that enhances sample quality by amplifying the contrast between conditional and unconditional (or negatively conditioned) likelihoods. This mechanism guides generation toward prompt-aligned outputs by explicitly pushing samples away from the lower-quality unconditional distribution. However, a challenge arises with many preference optimization methods, such as Diffusion Direct Preference Optimization (Diff-DPO). These methods often apply the same optimization objective to both the conditional and unconditional branches. This uniform reinforcement of preferred attributes, while shifting the overall distribution in the desired direction, fails to heighten the critical contrast that CFG relies on. Consequently, the model's ability to generate outputs aligned with user intent remains limited.

Diffusion Negative Preference Optimization (Diff-NPO) (Wang et al., 2025) was introduced to address this limitation. It enhances the CFG contrast by training two separate models: a standard "positive" model on preference data and a "negative" model on inverted preferences. When sampling, the negative model replaces the unconditional branch, explicitly steering generation away from undesirable attributes. While effective in principle, this dual-model approach suffers from two critical flaws. First, it doubles the training and inference costs, a prohibitive burden for large-scale models. Second, to enable effective sampling, the separately trained weights must be merged, which

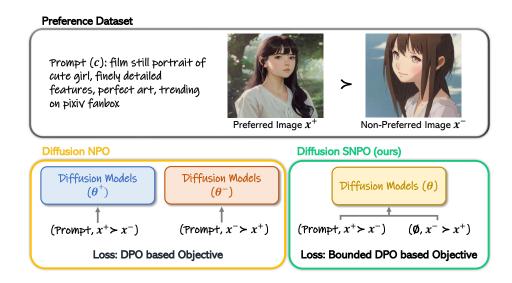


Figure 1: **Illustration of Diff-SNPO.** Winning and losing samples are denoted by x^+ and x^- , respectively. Unlike Diff-NPO, which trains two separate models for positive and negative preferences, Diff-SNPO integrates both into a single model. To ensure stable optimization in this unified setting, we introduce a Bounded DPO-based diffusion objective.

improves generation quality but weakens alignment. This merging introduces a distribution shift that attenuates the negative signal, reducing the model's ability to steer away from dispreferred outputs.

In this work, we ask: Can we achieve the benefits of explicit negative modeling without the costs and compromises of a dual-model architecture? We introduce Diffusion Simple Negative Preference Optimization (Diffusion-SNPO), a framework that integrates positive and negative preference signals into a single, unified network. Our approach leverages the inherent dual-branch structure of CFG-enabled models, training the conditional branch on positive preferences and the unconditional branch on negative (inverted) preferences. This design eliminates the need for separate models and weight merging, preserving a strong, explicit contrast between preferred and dispreferred distributions within one efficient architecture.

However, we find that naively applying this strategy results in an unintended side effect: the generated images become progressively blurrier as training continues. We attribute this to the likelihood decrease observed in the DPO algorithm (Rafailov et al., 2024; Pal et al., 2024; Cho et al., 2025), which, when coupled with conflicting gradients from flipped preferences between the conditional and unconditional branches, causes instability. As a result, the model converges to a suboptimal solution, with blurring becoming more pronounced as the likelihood of winning samples decreases. To address this, we adapt Bounded DPO (Cho et al., 2025), a preference optimization method designed to increase the likelihood of winning samples during training, to Diffusion Models. This adaptation stabilizes training by preventing the loss from being dominated by low-likelihood "losing" samples, thereby boosting the likelihood of winning samples throughout preference optimization.

In detail, our contributions can be summarized as follows:

- We identify the challenges of dual-model negative preference optimization (NPO), namely its high computational cost and the performance trade-offs inherent in its two separate model design.
- We demonstrate that simply applying opposing preferences on different branches within a single model leads to poor interactions with DPO-based algorithms, resulting in blurry images as training progresses.
- We introduce Diff-SNPO, a single-model negative preference optimization method for Diffusion Models. By adapting the BDPO algorithm and deriving an upper bound, our approach effectively eliminates the blurring effect caused by DPO in a negative preference optimization setting.

• We conduct extensive evaluations on the Pick-a-Pic v2 benchmark with both SD1.5 and SDXL, showing that **Diff-SNPO** delivers strong performance across multiple alignment metrics, while being more efficient in both training and inference.

2 Related Work

2.1 Preference Optimization

Preference optimization aims to align generative models with human judgments. One common strategy is to train a *reward model* that scores prompt–image pairs based on semantic or aesthetic quality, and then fine-tune the diffusion model to maximize these scores (Xu et al., 2023; Clark et al., 2024). Another direction draws on *policy optimization*, which frames denoising as a sequential decision process: DDPO (Black et al., 2024) applies reinforcement learning across the sampling steps, while DPOK (Fan et al., 2023) introduces a KL-regularized reward objective. More recently, methods that bypass explicit reward modeling by learning directly from curated positives or pairwise preferences have attracted growing interest (Wallace et al., 2024; Lu et al., 2025; Hong et al., 2025). While these approaches simplify training, they typically handle negative feedback only implicitly, by training models to favor preferred outputs in relative comparisons. This indirect treatment leaves little control over explicitly suppressing undesirable generations. In contrast, explicitly modeling dispreferred outcomes offers a more direct form of control: it steers generation away from undesirable regions of the distribution, reducing artifacts and improving alignment.

2.2 Negative Preference Optimization

In response to the limitations of standard preference optimization algorithms, recent work has explored dual-model strategies that explicitly separate positive and negative preferences. Diff-NPO (Wang et al., 2025) trains a negative model on inverted preferences and substitutes it for the unconditional branch during inference, which improves alignment but introduces a mismatch between training and inference due to weight interpolation. CHATS (Fu et al., 2025) addresses this issue by jointly training positive and negative models within a contrastive objective, achieving stronger alignment but at the cost of increased computation and memory from maintaining two networks. These challenges highlight the need for approaches that unify positive and negative preference modeling within a single framework. Our method, DIFF-SNPO, takes this direction by integrating both signals into a shared training objective, reducing overhead while better balancing alignment across positive and negative preferences.

3 PRELIMINARIES

3.1 DIFFUSION MODELS

Denoising Diffusion Probabilistic Models (DDPMs) (Ho et al., 2020) consist of a forward noising process and a learned reverse denoising process. The forward process gradually perturbs clean data x_0 with Gaussian noise according to a variance schedule $\{\beta_t\}_{t=1}^T$:

$$q(x_{1:T} \mid x_0) = \prod_{t=1}^{T} q(x_t \mid x_{t-1}), \quad q(x_t \mid x_{t-1}) = \mathcal{N}\left(x_t; \sqrt{1 - \beta_t} \, x_{t-1}, \beta_t I\right). \tag{1}$$

To invert this process, a neural network ϵ_{θ} parameterizes the reverse transitions by predicting the injected noise:

$$p_{\theta}(x_{t-1} \mid x_t) = \mathcal{N}\left(x_{t-1}; \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sqrt{1-\bar{\alpha}_t}} \epsilon_{\theta}(x_t, t)\right), \sigma_t^2 I\right), \tag{2}$$

where $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$, and σ_t^2 denotes the variance.

Model training minimizes a variational bound, which simplifies to a weighted noise-prediction loss:

$$\mathcal{L}_{\text{DDPM}} = \mathbb{E}_{x_0, \epsilon, t} \left[\omega(t) \| \epsilon - \epsilon_{\theta}(x_t, t) \|^2 \right], \tag{3}$$

with $\epsilon \sim \mathcal{N}(0, I)$ and $t \sim \mathcal{U}\{1, \dots, T\}$. The weighting function $\omega(t)$ controls the relative contribution of different timesteps, reflecting the varying difficulty of denoising across the diffusion

trajectory. This objective provides a simple and stable training criterion that underpins most modern diffusion-based generative models.

3.2 DIFFUSION DIRECT PREFERENCE OPTIMIZATION

Preference alignment is often formalized using the Bradley–Terry (BT) model (Kendall & Smith, 1940). Building on this idea, Direct Preference Optimization (DPO) (Rafailov et al., 2023) sidesteps the need to fit an explicit reward model by defining an implicit reward through likelihood ratios between the policy π_{θ} and a reference policy π_{ref} :

$$\mathcal{L}_{\text{DPO}}(\theta) = -\mathbb{E}_{(x_0^w, x_0^l, c) \sim \mathcal{D}} \Big[\log \sigma \Big(\beta \Big(\log \frac{\pi_{\theta}(x_0^w \mid c)}{\pi_{\text{ref}}(x_0^w \mid c)} - \log \frac{\pi_{\theta}(x_0^l \mid c)}{\pi_{\text{ref}}(x_0^l \mid c)} \Big) \Big) \Big]. \tag{4}$$

Here, $\beta>0$ is a temperature parameter that controls the strength of the preference signal, and $\sigma(\cdot)$ denotes the sigmoid function. Extending this objective to diffusion models requires assigning preferences over entire trajectories $x_{0:T}$. In this context, Diff-DPO (Wallace et al., 2024) derives an upper bound on the exact Diff-DPO objective:

$$\mathcal{L}_{\text{Diff-DPO}}(\theta; y) = -\mathbb{E}_{(x_0^w, x_0^l, c) \sim \mathcal{D}, t \sim \mathcal{U}[1, T]} \Big[\log \sigma \big(y T \omega(t) \beta (\Delta_t^w(c) - \Delta_t^l(c)) \big) \Big], \tag{5}$$

with

$$\Delta_t^w(c) = \left\| \epsilon^w - \epsilon_\theta(x_t^w, t, c) \right\|_2^2 - \left\| \epsilon^w - \epsilon_{\text{ref}}(x_t^w, t, c) \right\|_2^2, \tag{6}$$

$$\Delta_t^l(c) = \left\| \epsilon^l - \epsilon_\theta(x_t^l, t, c) \right\|_2^2 - \left\| \epsilon^l - \epsilon_{\text{ref}}(x_t^l, t, c) \right\|_2^2, \tag{7}$$

Here, x_t^w and x_t^l are noisy states obtained by applying the forward diffusion process to the clean samples x_0^w, x_0^l with corresponding noise terms $\epsilon^w, \epsilon^l \sim \mathcal{N}(0, I)$. The label $y \in \{+1, -1\}$ encodes the preference direction and $\omega(t)$ is a time-dependent weighting function.

3.3 DIFFUSION NEGATIVE PREFERENCE OPTIMIZATION

Diffusion Negative Preference Optimization (Diff-NPO) (Wang et al., 2025) extends standard preference learning by explicitly modeling undesirable behavior. It does so by training on an *inverted* preference dataset, where the roles of winners and losers are swapped. This yields a negatively aligned model, θ^- , which learns to assign higher likelihood to dispreferred samples. During sampling, Diff-NPO replaces the standard unconditional branch in classifier-free guidance with this negative model, leading to the following guidance formulation:

$$\epsilon_{\text{NPO}}(x_t, c) = \epsilon_{\theta^+}(x_t, c) + \omega \left[\epsilon_{\theta^+}(x_t, c) - \epsilon_{\theta^-}(x_t, \varnothing) \right], \tag{8}$$

where ϵ_{θ^+} is the positively aligned model trained on standard preference data, and ω is a guidance strength hyperparameter.

In practice, however, the positive and negative models that are trained independently often exhibit poor correlation, which undermines the effectiveness of ϵ_{θ^-} as a meaningful contrastive signal during sampling. To mitigate this, Diff-NPO applies a weight-merging procedure that combines the reference model $\theta_{\rm ref}$, the positive model θ^+ , and the negative model θ^- :

$$\hat{\theta}^{-} = \theta_{\text{ref}} + \alpha(\theta^{+} - \theta_{\text{ref}}) + \beta(\theta^{-} - \theta_{\text{ref}}), \tag{9}$$

where α and β are interpolation coefficients. In practice, the merged parameters $\hat{\theta}^-$ replace ϵ_{θ^-} in the guidance formulation of Eq. 9, yielding improved generation quality. However, this benefit comes at the cost of higher training overhead and weakened negative alignment.

4 METHOD

4.1 Limitations of Negative Preference Optimization

Dual-model approaches to Negative Preference Optimization (NPO) (Wang et al., 2025) offer direct way to model positive and negative preferences, training separate models— θ^+ and θ^- —for each.

However, this design introduces significant practical overhead. Maintaining two networks doubles memory and compute requirements, increases training time. These costs quickly escalate with large diffusion backbones, making dual-model setups difficult to scale in practice.

At inference time, using the two models separately preserves the intended contrast between positive and negative preferences. However, this approach often compromises sample quality, especially when generating from the negatively aligned model θ^- . To mitigate this, Diff-NPO introduces a merged parameter formulation, $\hat{\theta}^-$, which interpolates between θ^+ , θ^- , and the reference backbone $\theta_{\rm ref}$ as shown in Eq. 9. This interpolation enhances generation quality and increases agreement between the two models' predictions (Wang et al., 2025), but it also biases the merged model toward θ^+ , reducing the influence of θ^- . As illustrated in Figure 2, this trade-off between fidelity and alignment exposes a deeper limitation of the dual-model paradigm: its reliance on separate networks and post-hoc merging undermines the very preference contrast that NPO is designed to enforce. To overcome these issues, we explore single-model formulations that unify both positive and negative preference signals within a shared architecture—offering a more robust and scalable solution, which we introduce in the next section.

4.2 ISSUES WITH NAIVE DIFF-SNPO

Standard diffusion models inherently support both conditional and unconditional branches within a single network—a design popularized by Classifier-Free Guidance (Ho & Salimans, 2021). This built-in dual-branch structure presents a natural alternative to the dual-model Diff-NPO setup: rather than maintaining two separate networks, one can leverage the existing architecture by updating the conditional branch with preferred samples and the unconditional (or negatively conditioned) branch with dispreferred ones. This approach preserves the contrast between preference signals while avoiding the redundancy and overhead of training and managing separate models.

Specifically, let $Y \in \{+1, -1\}$ be a branch label with $\Pr(Y = +1) = 1 - p$ and $\Pr(Y = -1) = p$, where $p \in [0, 1]$ is the CFG dropout probability. The effective conditioning is

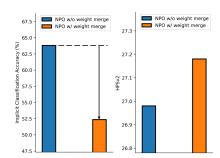


Figure 2: Negative implicit classification accuracy (left) and HPSv2 score (right) on SD1.5. Weight merging leads to a substantial drop in implicit accuracy while increasing the reward score, highlighting a trade-off between the two.

$$\tilde{c}(Y) = \begin{cases} c, & Y = +1 \\ \varnothing, & Y = -1 \end{cases}$$
 (conditional branch),

where c denotes the conditioning input (e.g., a text prompt) and \varnothing indicates null conditioning, as in CFG. A Naive Diff-SNPO objective can thus be constructed by adapting the Diff-DPO objective in Eq. 5:

$$\mathcal{L}_{\text{Naive Diff-SNPO}}(\theta) = -\mathbb{E}_{(x_0^w, x_0^l) \sim \mathcal{D}, t, Y} \left[\log \sigma (YT \,\omega(t) \,\beta \, \left(\Delta_t^w(\tilde{c}(Y)) - \Delta_t^l(\tilde{c}(Y)) \right) \right) \right], \tag{10}$$

While Naive Diff-SNPO offers a simple and intuitive way to model flipped preferences within a single network, it suffers from significant degradation in generation quality over training. As illustrated in Fig. 3, this naive objective yields progressively blurrier generations over training, with reduced contrast and attenuated high-frequency detail. We attribute this effect to a well-documented property of the DPO objective: improvements in the pairwise margin often result in a reduction in likelihood for both candidates, penalizing the loser more heavily, rather than a reliable increase in the likelihood of the winner (Pal et al., 2024; Rafailov et al., 2024; Cho et al., 2025).

To quantify this effect on preferred samples, we track the win-sample likelihood ratio against a fixed reference model, given by:

$$\frac{\pi_{\theta}(x^w)}{\pi_{\text{ref}}(x^w)} \approx \mathbb{E}_{x_0^w, t} \left[e^{\Delta_t^w(c)} \right].$$

Under Naive Diff-SNPO, we observe a consistent decrease in the relative win probability of preferred samples over the course of training (Fig. 3), indicating that the model suppresses likelihoods

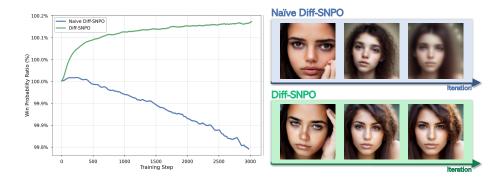


Figure 3: Win probability ratio (left) and generated samples across training iterations (right). Naive-SNPO exhibits a steady decline in win probability ratio, resulting in progressively blurred outputs, whereas Diff-SNPO preserves stability and yields increasingly preference-aligned images.

across both branches rather than reinforcing the preferred ones. This undesirable trend, combined with the structure of the optimization signals, leads to a characteristic blurring effect. Specifically, Naive Diff-SNPO applies symmetric but opposing updates to the conditional and unconditional branches—encouraging one to increase the likelihood of a sample while the other decreases it. Since both branches share model parameters, these conflicting gradients interfere with each other, and the model resolves this conflict by averaging them. In generative settings, this averaging dampens contrast and degrades fine detail, resulting in blurry outputs. To overcome this limitation, we seek to break the destructive symmetry in the update rule. In the next section, we introduce an asymmetric preference optimization approach that biases learning toward increasing win probability, thereby mitigating gradient interference and addressing the blurring effect.

4.3 DIFF-SNPO

Recently, Cho et al. (2025) introduced Bounded DPO (BDPO) to address a key shortcoming of standard DPO. As the model reduces the probability assigned to the losing sample, the corresponding lose sample log-likelihood term, $\log \pi_{\theta}(\mathbf{y}_l \mid \mathbf{x})$, in the objective grows disproportionately large, causing the loss to become dominated by the loser. This skews the gradient signal and can even drive updates that decrease the likelihood of the preferred (winning) output, despite the preference label.

To mitigate this issue, BDPO replaces the losing term with a mixture distribution that includes a non-vanishing contribution from the reference policy:

$$\pi_{\text{mix}}(\mathbf{y} \mid \mathbf{x}) = \lambda \, \pi_{\theta}(\mathbf{y} \mid \mathbf{x}) + (1 - \lambda) \, \pi_{\text{ref}}(\mathbf{y} \mid \mathbf{x}), \qquad \lambda \in (0, 1),$$

which leads to the modified objective

$$\mathcal{L}_{\text{BDPO}}(\pi_{\theta}; \pi_{\text{ref}}) = -\mathbb{E}_{(x_0^w, x_0^l, c) \sim \mathcal{D}} \left[\log \sigma \left(\beta \left[\log \frac{\pi_{\theta}(x_0^w|c)}{\pi_{\text{ref}}(x_0^w|c)} - \log \frac{\pi_{\text{mix}}(x_0^l|c)}{\pi_{\text{ref}}(x_0^l|c)} \right] \right) \right]. \tag{11}$$

This modification bounds the contribution of the loser term and prevents it from overwhelming the loss, thereby preserving the intended effect of preference optimization: promoting the winning sample. In addition, BDPO retains the same global minimizers as DPO while enforcing a lower bound on the winning likelihood, offering stronger stability guarantees throughout training (Cho et al., 2025).

To adapt BDPO to diffusion models, similar to Diff-DPO, we define the trajectory-level reward

$$r_{\theta}(c, \mathbf{x}_0) = \mathbb{E}_{\mathbf{x}_{1:T} \sim p_{\theta}(\cdot | \mathbf{x}_0, c)} \Big[R(c, \mathbf{x}_{0:T}) \Big]. \tag{12}$$

The corresponding Diffusion BDPO objective is

$$\mathcal{L}_{\text{Diff-BDPO}}(\theta) = -\mathbb{E}_{(\mathbf{x}_{0:T}^w, \mathbf{x}_{0:T}^l, c) \sim \mathcal{D}} \left[\log \sigma \left(\beta \left[\log \frac{p_{\theta}(\mathbf{x}_{0:T}^w | c)}{p_{\text{ref}}(\mathbf{x}_{0:T}^w | c)} - \log \frac{\pi_{\text{mix}}(\mathbf{x}_{0:T}^l | c)}{p_{\text{ref}}(\mathbf{x}_{0:T}^l | c)} \right] \right) \right]. \tag{13}$$

Following Diff-DPO, we upper-bound this objective via its ELBO and then apply Jensen's inequality to obtain the result below (see Appendix A.2 for the full derivation):

$$\mathcal{L}_{\text{Diff-BDPO-UB}}(\theta) = -\mathbb{E}_{(x_0^w, x_t^l, c) \sim \mathcal{D}, t \sim \mathcal{U}[1, T]} \left[\log \sigma \left(-\beta \left(m(x_t^w, c) - m_{\text{mix}}(x_t^l, c) \right) \right) \right]. \tag{14}$$

Here the per-step terms are:

$$d_{\theta}(x_{t}, \epsilon, t, c) = T \omega(t) \|\epsilon - \epsilon_{\theta}(x_{t}, t, c)\|_{2}^{2}, \quad d_{\text{ref}}(x_{t}, \epsilon, t, c) = T \omega(t) \|\epsilon - \epsilon_{\text{ref}}(x_{t}, t, c)\|_{2}^{2}, \quad (15)$$

$$m(x_{t}, c) = d_{\theta}(x_{t}, \epsilon, t, c) - d_{\text{ref}}(x_{t}, \epsilon, t, c), \quad (16)$$

$$m_{\text{mix}}(x_t, c) = -\log\left(\lambda e^{-d_{\theta}(x_t, \epsilon, t, c)} + (1 - \lambda) e^{-d_{\text{ref}}(x_t, \epsilon, t, c)}\right) - d_{\text{ref}}(x_t, \epsilon, t, c). \tag{17}$$

where $\mathcal{L}_{\text{Diff-BDPO-UB}}$ denotes the upper bound approximation of Diff-BDPO.

Building on this, we define our final Diff-SNPO objective, which applies our introduced Diff-BDPO-UB objective into a single model negative preference optimization framework:

$$\mathcal{L}_{\text{SNPO}}(\theta) = -\mathbb{E}_{(\mathbf{x}^w, \mathbf{x}^l, c) \sim \mathcal{D}, t \sim p(t), Y} \left[\log \sigma \left(\beta \left(m(\tilde{x}^w(Y), \tilde{c}(Y)) - m_{\text{mix}}(\tilde{x}^l(Y), \tilde{c}(Y)) \right) \right) \right]. \tag{18}$$

where $\tilde{x}^{w/l}(Y)$ is given as:

$$\tilde{x}^{w/l}(Y) = \begin{cases} x^{w/l}, & \text{if } Y = +1 & \text{(conditional branch)} \\ x^{l/w}, & \text{if } Y = -1 & \text{(unconditional/null branch)}. \end{cases}$$

As shown in Fig. 3, Diff-SNPO avoids the decline in winner likelihood observed with Naive Diff-SNPO: the estimated log-likelihood of winning samples steadily improves throughout training. In line with this, Fig. 3 further shows that Diff-SNPO prevents the progressive low-contrast "blurring" artifact characteristic of the naive training objective, effectively addressing its key shortcoming.

5 EXPERIMENTS

5.1 EXPERIMENTAL SETUP

Datasets and Models. We fine-tune both Stable Diffusion v1.5 (SD1.5) (Rombach et al., 2022) and Stable Diffusion XL (SDXL) (Podell et al., 2024) using our proposed Diffusion SNPO objective in Eq. 18. For training, we utilize the Pick-a-Pic v2 (Kirstain et al., 2023) corpus, a large-scale human preference dataset comprising 851,293 image pairs across 58,960 unique prompts.

Training Details. Models are initialized from the publicly available SD1.5 (Rombach et al., 2022) (CreativeML Open RAIL-M license) and SDXL (Podell et al., 2024) (MIT license) checkpoints. Training is conducted using the AdamW (Loshchilov & Hutter, 2019) optimizer with a learning rate of 2.048×10^{-8} . We use a batch size of 512 and train for 3,000 steps on SD1.5, and a larger batch size of 2,048 for 625 steps on SDXL. All experiments are run on $8 \times \text{NVIDIA A6000 GPUs}$ using distributed data parallelism. Total training time is approximately 12 hours for SD1.5 and 68 hours for SDXL. The regularization coefficient is set to $\beta = 2000$ for SD1.5 and $\beta = 5000$ for SDXL. Unless stated otherwise, Diff-SNPO is trained with an interpolation parameter of $\lambda = 0.9$.

Baselines. We compare Diff-SNPO against both negative preference optimization methods and standard alignment baselines. Specifically, we evaluate Diff-NPO (Wang et al., 2025), CHATS (Fu et al., 2025), and Diff-DPO (Wallace et al., 2024), along with the original pretrained models for SD1.5 (Rombach et al., 2022) and SDXL (Podell et al., 2024).

Evaluation Protocol. All models are evaluated using the DDIM (Song et al., 2021) sampler with 50 inference steps and a classifier-free guidance scale of 7.5. For Diff-NPO and CHATS, we adopt the hyperparameters specified in their released code, as these methods modify the sampling pipeline. Performance is assessed using five widely adopted reward models: PickScore (Kirstain et al., 2023), HPSv2 (Wu et al., 2023), ImageReward (Xu et al., 2023), and Aesthetics Score (Schuhmann, 2023). To account for stochasticity, we report average results across four random seeds, with each model generating 2,000 images using prompts from the Pick-a-Pic v2 test set.

Table 1: Comparison of Diff-SNPO with baseline methods on SD1.5 and SDXL backbones using Pick-a-Pic v2. Diff-SNPO consistently achieves the highest scores across most human preference metrics, reflecting improved alignment and visual quality. For clarity, the best-performing method in each metric is shown in **bold**, and the second-best is underlined.

Model	Method	HPSv2	Pick Score	Aesthetic Score	Image Reward
SD1.5	Baseline	26.24	20.64	5.2849	0.1221
	Diff-DPO (Wallace et al., 2024)	26.55	21.01	5.3823	0.2968
	Diff-NPO (Wang et al., 2025)	26.92	21.46	5.5381	0.3786
	CHATS (Fu et al., 2025)	27.20	21.05	5.6845	0.2995
	Diff-BDPO (Ours)	26.64	21.15	5.4457	0.3165
	Diff-SNPO (Ours)	27.23	22.24	5.6258	0.6936
SDXL	Baseline	27.43	22.13	5.8850	0.7605
	Diff-DPO (Wallace et al., 2024)	28.09	22.59	5.8884	0.9841
	Diff-NPO (Wang et al., 2025)	28.30	22.67	5.9449	0.9847
	CHATS (Fu et al., 2025)	28.25	22.34	5.8785	1.0543
	Diff-BDPO (Ours)	28.13	22.36	5.8037	0.9946
	Diff-SNPO (Ours)	28.33	22.69	5.8129	<u>1.0100</u>

5.2 QUANTITATIVE RESULTS

Table 1 presents quantitative comparisons of Diff-SNPO against established baselines on SD1.5 and SDXL backbones.

On SD1.5, Diff-SNPO consistently outperforms prior methods across most human preference metrics. Notably, it achieves 27.23 on HPSv2 and 22.24 on PickScore, marking a substantial improvement over both Diff-DPO and the more stable Diff-BDPO. This indicates that the gains stem not just from BDPO's stabilization, but from negative preference optimization itself. While CHATS attains a slightly higher Aesthetic Score, it lags behind on all other metrics, suggesting that its more visually pleasing outputs come at the cost of semantic alignment. Diff-SNPO also surpasses its dual-model counterpart, Diff-NPO, across all metrics while using only half the computation. We attribute this to its single-model design, which preserves negative preference alignment better than Diff-NPO, where accuracy drops notably after model merging.

On SDXL, Diff-SNPO delivers strong and competitive results, with HPSv2 and ImageReward scores of 28.33 and 1.01, comparable to state-of-the-art methods. Its Aesthetic Score (5.81) is slightly below DPO (5.89), and the large advantage it shows over Diff-NPO on SD1.5 becomes smaller on SDXL. This stems from Diff-SNPO's bias toward optimizing "preferred" samples in the dataset. Although these samples are favored over their losing counterparts, they have been shown to be less aesthetically pleasing than the already strong outputs of the base SDXL model(). As a result, focusing on them may slightly reduce aesthetic quality and limit gains on stronger backbones. Similar trends appear with CHATS and Diff-BDPO, suggesting the limitation comes from the dataset rather than the alignment method itself. Even so, Diff-SNPO retains important practical benefits: it requires only half the computation of Diff-NPO and, with its single-model architecture, also enables faster sampling—making it both efficient and scalable.

5.3 Impact of the Mixing Parameter λ on Diff-SNPO

Table 2 reports results on SD1.5, with all hyperparameters fixed as in the main experiments except for the mixing parameter λ , show that performance remains stable across different λ values. In particular, $\lambda=0.9$ as it consistently delivers strong performance. In contrast, the $\lambda=1.0$ setting, which corresponds to Naive-SNPO, shows a notable drop across both reward metrics. As discussed in Section 4.2, this decline results from blurring artifacts that arise during training when the win likelihood decreases during single model Negative Prefer-

Table 2: **Effect of varying** λ **on Diff-SNPO** (**SD1.5**). Performance remains stable across different λ values, while Naive Diff-SNPO ($\lambda = 1.0$) suffers a pronounced drop, reflecting the onset of severe blurring artifacts.

λ	HPSv2↑	Aesthetic ↑
0.1	27.17	5.5818
0.5	27.27	5.5893
0.9	27.23	5.6258
1.0	25.84	5.3169

Table 3: **Training cost comparison.** Experiments were conducted on a single A6000 GPU with a batch size of 8. Dual-model approaches require more memory and result in slower relative speed The best result in each column is highlighted in **bold**.

Method	Memory (GB) ↓	Time / Step (s) \downarrow	Relative Speed ↑
Diff-NPO	41.4×2	12.12×2	1.00×
CHATS	45.3	13.08	1.85×
Diff-SNPO (Ours)	41.4	12.12	$2.00 \times$

ence Optimization. These findings suggest that $\lambda=0.9$ serves as a favorable operating point, effectively balancing the stability benefits of BDPO with a reduced bias toward win probability. This balance results in more robust training dynamics and improved empirical performance.

5.4 Comparing Negative Preference Alignment Across Methods

From Table 4, we observe that both Diff-DPO and Diff-BDPO achieve similar negative implicit accuracy in their unconditional branches, consistently falling below 50%. This outcome aligns with their training setup, as the unconditional branches are optimized using positive preference alignment, which limits their capacity to capture negative preference signals. In contrast to Diff-DPO and Diff-BDPO, Diff-NPO exhibits a notable decline in negative implicit accuracy after weight merging, with a drop of over 10 %, falling

Table 4: **Negative preference implicit classification accuracy** Parentheses denote Diff-NPO without weight merging. Diff-SNPO achieves higher negative implicit accuracy than Diff-NPO improving its negative alignment.

Method	Neg. Implicit Acc. (%)		
Diff-DPO	31.86		
Diff-BDPO	32.04		
Diff-NPO	52.34 (63.80)		
Diff-SNPO	57.45		

below Diff-SNPO and approximately 5% behind Diff-NPO. These findings suggest that Diff-SNPO offers a more effective approach for preserving negative alignment in the unconditional branch.

5.5 Training Computational Cost

Beyond alignment performance, the practicality of a preference optimization algorithm also depends on its computational efficiency. To assess this, we compare the training computational cost of dual-model approaches (CHATS and Diff-NPO) against our single-model Diff-SNPO. As reported in Table 3, the single-model design of Diff-SNPO yields substantial efficiency gains in both memory usage and training time. By eliminating the need to train two separate networks, Diff-SNPO reduces memory consumption and achieves a 2× speedup in per-step training time relative to its dual-model counterpart, Diff-NPO. A similar advantage is observed over CHATS: while both methods require comparable memory, Diff-SNPO trains faster because it optimizes the conditional and unconditional branches in parallel, whereas CHATS processes them sequentially. In summary, Diff-SNPO combines lower memory overhead with faster training, establishing it as a more efficient and scalable alternative for negative preference optimization. A detailed comparison of inference cost, which further highlights the efficiency of our single-model approach, can be found in Appendix A.4.

6 Conclusion

In conclusion, we propose Diff-SNPO, a single-model framework for Negative Preference Optimization that achieves strong performance while simplifying the training pipeline. In contrast to prior approaches that require two separate models, Diff-SNPO unifies conditional and unconditional branches within a single architecture, thereby eliminating redundant computation and substantially improving efficiency in both training and inference. Our experiments demonstrate that this streamlined design not only matches, but often exceeds the performance of existing methods, highlighting its effectiveness in preserving negative preference alignment. Beyond raw performance, the reduced computational footprint and faster sampling make Diff-SNPO a practical and scalable solution, lowering the barriers to applying preference optimization in real-world generative modeling tasks.

REFERENCES

- Kevin Black, Michael Janner, Yilun Du, Ilya Kostrikov, and Sergey Levine. Training diffusion models with reinforcement learning. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=YCWjhGrJFD.
- Jay Hyeon Cho, JunHyeok Oh, Myunsoo Kim, and Byung-Jun Lee. Rethinking dpo: The role of rejected responses in preference misalignment, 2025. URL https://arxiv.org/abs/2506.12725.
- Kevin Clark, Paul Vicol, Kevin Swersky, and David J. Fleet. Directly fine-tuning diffusion models on differentiable rewards. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=1vmSEVL19f.
- Ying Fan, Olivia Watkins, Yuqing Du, Hao Liu, Moonkyung Ryu, Craig Boutilier, Pieter Abbeel, Mohammad Ghavamzadeh, Kangwook Lee, and Kimin Lee. Reinforcement learning for fine-tuning text-to-image diffusion models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=80TPepXzeh.
- Minghao Fu, Guo-Hua Wang, Liangfu Cao, Qing-Guo Chen, Zhao Xu, Weihua Luo, and Kaifu Zhang. CHATS: Combining human-aligned optimization and test-time sampling for text-to-image generation. In Forty-second International Conference on Machine Learning, 2025. URL https://openreview.net/forum?id=D4Y71nbGRg.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. URL https://openreview.net/forum?id=qw8AKxfYbI.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 6840–6851. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf.
- Jonathan Ho, Tim Salimans, Alexey A. Gritsenko, William Chan, Mohammad Norouzi, and David J. Fleet. Video diffusion models. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL https://openreview.net/forum?id=f3zNgKga_ep.
- Jiwoo Hong, Sayak Paul, Noah Lee, Kashif Rasul, James Thorne, and Jongheon Jeong. Margin-aware preference optimization for aligning diffusion models without reference. In *First Workshop on Scalable Optimization for Efficient and Adaptive Foundation Models*, 2025. URL https://openreview.net/forum?id=JMVWaMR53s.
- M. G. Kendall and B. Babington Smith. On the method of paired comparisons. *Biometrika*, 31 (3-4):324–345, 03 1940. ISSN 0006-3444. doi: 10.1093/biomet/31.3-4.324. URL https://doi.org/10.1093/biomet/31.3-4.324.
- Yuval Kirstain, Adam Polyak, Uriel Singer, Shahbuland Matiana, Joe Penna, and Omer Levy. Picka-pic: An open dataset of user preferences for text-to-image generation. 2023.
- Shufan Li, Konstantinos Kallidromitis, Akash Gokul, Yusuke Kato, and Kazuki Kozuka. Aligning diffusion models by optimizing human utility. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. URL https://openreview.net/forum?id=MTMShU5QaC.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=Bkg6RiCqY7.
 - Yunhong Lu, Qichao Wang, Hengyuan Cao, Xierui Wang, Xiaoyin Xu, and Min Zhang. Inpo: Inversion preference optimization with reparametrized ddim for efficient diffusion model alignment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 28629–28639, June 2025.

- Arka Pal, Deep Karkhanis, Samuel Dooley, Manley Roberts, Siddartha Naidu, and Colin White. Smaug: Fixing failure modes of preference optimisation with dpo-positive. *arXiv* preprint *arXiv*:2402.13228, 2024.
- Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: Improving latent diffusion models for high-resolution image synthesis. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=di52zR8xgf.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=HPuSIXJaa9.
- Rafael Rafailov, Joey Hejna, Ryan Park, and Chelsea Finn. From \$r\$ to \$q^*\$: Your language model is secretly a q-function. In *First Conference on Language Modeling*, 2024. URL https://openreview.net/forum?id=kEVcNxtqXk.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10684–10695, June 2022.
- Ludan Ruan, Yiyang Ma, Huan Yang, Huiguo He, Bei Liu, Jianlong Fu, Nicholas Jing Yuan, Qin Jin, and Baining Guo. Mm-diffusion: Learning multi-modal diffusion models for joint audio and video generation. In *CVPR*, 2023.
- Christoph Schuhmann. Laion-aesthetics. https://github.com/christophschuhmann/improved-aesthetic-predictor, October 2023.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=StlgiarCHLP.
- Bram Wallace, Meihua Dang, Rafael Rafailov, Linqi Zhou, Aaron Lou, Senthil Purushwalkam, Stefano Ermon, Caiming Xiong, Shafiq Joty, and Nikhil Naik. Diffusion model alignment using direct preference optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8228–8238, June 2024.
- Fu-Yun Wang, Yunhao Shui, Jingtan Piao, Keqiang Sun, and Hongsheng Li. Diffusion-NPO: Negative preference optimization for better preference aligned generation of diffusion models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=iJi7nz5Cxc.
- Xiaoshi Wu, Yiming Hao, Keqiang Sun, Yixiong Chen, Feng Zhu, Rui Zhao, and Hongsheng Li. Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. *arXiv* preprint arXiv:2306.09341, 2023.
- Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. Imagereward: Learning and evaluating human preferences for text-to-image generation. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=JVzeOYEx6d.

A APPENDIX

A.1 THE USE OF LLMS

We used LLMs solely for light editing such as correcting grammatical errors and polishing some words. They did not contribute to research ideation, experiments, analysis, or substantive writing.

A.2 DERIVATION OF THE DIFFUSION-BDPO UPPER BOUND

We begin by recalling the Diffusion-BDPO objective:

$$\mathcal{L}_{\text{Diff-BDPO}}(\theta) = -\mathbb{E}_{(\mathbf{x}_{0:T}^w, \mathbf{x}_{0:T}^l, c) \sim \mathcal{D}} \left[\log \sigma \left(\beta \left[\log \frac{p_{\theta}(\mathbf{x}_{0:T}^w \mid c)}{p_{\text{ref}}(\mathbf{x}_{0:T}^w \mid c)} - \log \frac{p_{\text{mix}}(\mathbf{x}_{0:T}^l \mid c)}{p_{\text{ref}}(\mathbf{x}_{0:T}^l \mid c)} \right] \right) \right], \quad (19)$$

where the mixture distribution is defined as

$$p_{\text{mix}}(\mathbf{x}_{0:T}^l \mid c) = \lambda p_{\theta}(\mathbf{x}_{0:T}^l \mid c) + (1 - \lambda) p_{\text{ref}}(\mathbf{x}_{0:T}^l \mid c), \qquad \lambda \in (0, 1).$$

$$(20)$$

To obtain a computationally tractable form, we first decompose $\log p_{\theta}(\mathbf{x}_{0:T} \mid c)$ using the backward process of diffusion models:

$$\log p_{\theta}(\mathbf{x}_{0:T} \mid c) = \log p_{\theta}(\mathbf{x}_T \mid c) + \log \prod_{t=1}^{T} p_{\theta}(\mathbf{x}_{t-1} \mid \mathbf{x}_t, c)$$
(21)

$$= \log p_{\theta}(\mathbf{x}_T \mid c) + \sum_{t=1}^{T} \log p_{\theta}(\mathbf{x}_{t-1} \mid \mathbf{x}_t, c)$$
 (22)

$$= C - \sum_{t=1}^{T} \omega(t) \|\epsilon_{\theta}(x_t, c, t) - \epsilon\|_{2}^{2},$$
 (23)

where C is a parameter-independent constant.

Next, consider the mixed log-probability:

$$\log p_{\text{mix}}(\mathbf{x}_{0:T}^{l} \mid c) = \log \left(\lambda p_{\theta}(\mathbf{x}_{0:T}^{l} \mid c) + (1 - \lambda) p_{\text{ref}}(\mathbf{x}_{0:T}^{l} \mid c) \right)$$

$$= C + \log \left(\lambda e^{\sum_{t=1}^{T} \log p_{\theta}(\mathbf{x}_{t-1}^{l} \mid \mathbf{x}_{t}, c)} + (1 - \lambda) e^{\sum_{t=1}^{T} \log p_{\text{ref}}(\mathbf{x}_{t-1}^{l} \mid \mathbf{x}_{t}, c)} \right). \tag{24}$$

For brevity, define

$$a_t := \log p_{\theta}(\mathbf{x}_{t-1}^l \mid \mathbf{x}_t, c), \qquad b_t := \log p_{\text{ref}}(\mathbf{x}_{t-1}^l \mid \mathbf{x}_t, c), \qquad t = 1, \dots, T.$$
 (25)

Then equation 24 becomes

$$\log p_{\min}(\mathbf{x}_{0:T}^{l} \mid c) = C + \log \left(\lambda e^{\sum_{t=1}^{T} a_t} + (1 - \lambda) e^{\sum_{t=1}^{T} b_t} \right).$$
 (26)

Applying the discrete n-factor Hölder inequality yields

$$\lambda e^{\sum_{t=1}^{T} a_t} + (1 - \lambda) e^{\sum_{t=1}^{T} b_t} \le \prod_{t=1}^{I} \left(\lambda e^{p_t a_t} + (1 - \lambda) e^{p_t b_t} \right)^{1/p_t}, \tag{27}$$

for exponents $p_1, \ldots, p_T \ge 1$ with $\sum_{t=1}^T 1/p_t = 1$.

Taking logs and substituting back gives

$$\log p_{\min}(\mathbf{x}_{0:T}^{l} \mid c) \le C + \sum_{t=1}^{T} \frac{1}{p_t} \log \left(\lambda e^{p_t a_t} + (1 - \lambda) e^{p_t b_t} \right). \tag{28}$$

Choosing $p_t = T$ for all t (so that $\sum_t 1/p_t = 1$) produces the uniform bound

$$\log p_{\min}(\mathbf{x}_{0:T}^l \mid c) \le C + \frac{1}{T} \sum_{t=1}^T \log \left(\lambda e^{T \log p_{\theta}(\mathbf{x}_{t-1}^l \mid \mathbf{x}_t, c)} + (1 - \lambda) e^{T \log p_{\text{ref}}(\mathbf{x}_{t-1}^l \mid \mathbf{x}_t, c)} \right). \tag{29}$$

Table 5: Comparison of Diff-SNPO with baseline methods on SD1.5 and SDXL backbones on HPDv2. Diff-SNPO achieves the best scores across most human preference metrics, indicating superior alignment and visual quality. For each metric, the top-performing method is **bolded**, while the second-best is underlined.

Model	Method	HPSv2	Pick Score	Aesthetic Score	Image Reward
SD1.5	Baseline	26.52	20.80	5.3406	-0.0521
	Diff-DPO (Wallace et al., 2024)	26.85	21.26	5.4716	0.1620
	Diff-NPO (Wang et al., 2025)	27.18	21.71	5.5842	0.2287
	CHATS (Fu et al., 2025)	27.68	21.45	5.8605	0.3699
	Diff-BDPO (Ours)	27.24	21.57	5.5677	0.3576
	Diff-SNPO (Ours)	27.85	22.86	<u>5.7737</u>	0.8093
SDXL	Baseline	28.01	22.80	5.9887	0.8781
	Diff-DPO (Wallace et al., 2024)	28.55	23.12	6.0296	1.0788
	Diff-NPO (Wang et al., 2025)	28.87	23.29	6.0816	1.1305
	CHATS (Fu et al., 2025)	28.82	22.85	5.9868	1.1165
	Diff-BDPO (Ours)	28.62	22.96	5.9562	1.0911
	Diff-SNPO (Ours)	28.89	23.34	5.9876	1.1460

Upper bound on the objective. Substituting equation 29 and equation 23 into equation 19, we obtain

$$\mathcal{L}_{\text{Diff-BDPO}}(\theta) \le -\mathbb{E}_{(\mathbf{x}_{0:T}^w, \mathbf{x}_{0:T}^l, c) \sim \mathcal{D}} \left[\log \sigma \left(\mathbb{E}_{t \sim \mathcal{U}[1, T]} \left[-\beta \, m(x_t^w, c) - m_{\text{mix}}(x_t^l, c) \right] \right) \right]$$
(30)

$$\leq -\mathbb{E}_{(x_0^w, x_0^l, c) \sim \mathcal{D}, t \sim \mathcal{U}[1, T]} \left[\log \sigma \left(-\beta \left(m(x_t^w, c) - m_{\text{mix}}(x_t^l, c) \right) \right) \right]$$

$$= \mathcal{L}_{\text{Diff-BDPO-UB}}(\theta), \qquad \text{(by Jensen's inequality)}. \tag{31}$$

Here the per-step terms are defined as

$$d_{\theta}(x_t, \epsilon, t, c) = T \omega(t) \|\epsilon - \epsilon_{\theta}(x_t, t, c)\|_2^2, \quad d_{\text{ref}}(x_t, \epsilon, t, c) = T \omega(t) \|\epsilon - \epsilon_{\text{ref}}(x_t, t, c)\|_2^2, \quad (32)$$

$$m(x_t, c) = d_{\theta}(x_t, \epsilon, t, c) - d_{\text{ref}}(x_t, \epsilon, t, c), \tag{33}$$

$$m_{\text{mix}}(x_t, c) = -\log\left(\lambda e^{-d_{\theta}(x_t, \epsilon, t, c)} + (1 - \lambda) e^{-d_{\text{ref}}(x_t, \epsilon, t, c)}\right) - d_{\text{ref}}(x_t, \epsilon, t, c). \tag{34}$$

In summary, by decomposing the diffusion likelihood and applying Hölder's inequality, we obtain a tractable upper bound on the original Diffusion-BDPO objective, expressed in terms of per-step denoising errors equation 32—equation 34.

A.3 ADDITIONAL QUANTITATIVE RESULTS

Table 5 presents the results on the HPDv2 test set, showing trends consistent with those observed on the Pick-a-Pic benchmark. On the SD1.5 backbone, Diff-SNPO delivers a clear improvement, raising HPSv2 to 22.86 and ImageReward to 0.81, underscoring its effectiveness in capturing human preferences. On the more advanced SDXL backbone, it continues to perform strongly, remaining competitive with the leading approach, Diff-NPO. Crucially, Diff-SNPO achieves this performance with significantly lower computational cost and faster sampling, benefits enabled by its single-network design. This efficiency advantage highlights its practicality, offering state-of-the-art alignment quality while maintaining scalability and runtime efficiency.

A.4 INFERENCE COMPUTATIONAL COST

Table 6 reports the inference throughput of Diff-SNPO and other negative preference optimization methods on SD1.5. The efficiency of the single-model design is particularly evident at inference: Diff-SNPO achieves the highest throughput among all methods, processing more than twice as many images per second as the dual-model baselines.

Table 6: **Inference cost comparison.** Single-image inference cost with SD1.5 on a single A6000 GPU.

$\mathbf{Method} \!\!\downarrow$	$\begin{tabular}{ll} Throughput (img/s) \uparrow \\ \end{tabular}$
CHATS	0.11
DiffNPO	0.27
Diff-SNPO (Ours)	0.48

By contrast, CHATS and Diff-NPO suffer from slower inference due to the need for two sequential forward passes—one for each model—combined with additional passes required by their sampling schemes. These results underscore Diff-SNPO's ability to perform fast, parallel inference without compromising alignment quality.

Overall, our findings demonstrate that Diff-SNPO effectively balances preference alignment with computational efficiency. Its single-model design reduces training cost while delivering faster inference, making it well-suited for large-scale training as well as real-time or resource-constrained deployment scenarios.

A.5 QUALITATIVE RESULTS

Figure 4 and Figure 5 present additional image samples generated by Diff-SNPO and various preference alignment methods on SD 1.5 and SDXL using prompts from HPDv2 test set respectively.

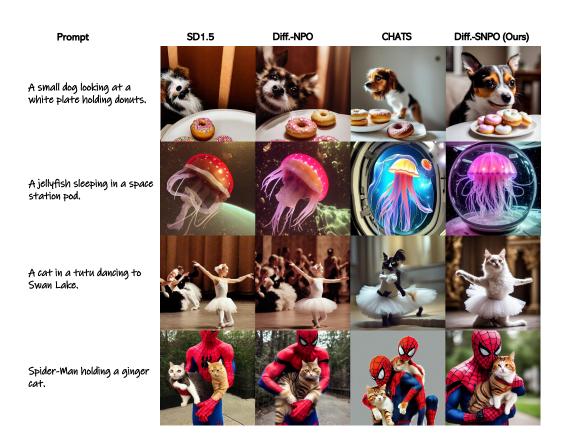


Figure 4: Side-by-side comparison of images generated by related methods on HPDv2 using SD1.5.

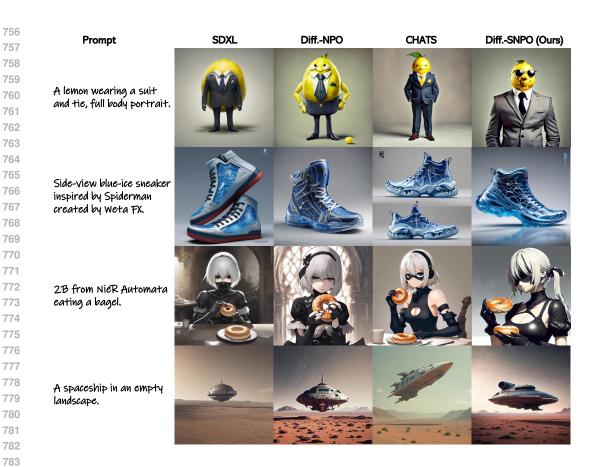


Figure 5: Side-by-side comparison of images generated by related methods on HPDv2 using SDXL.